

12/10/99
JC682 U.S. PTO

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER, L.L.P.

1300 I STREET, N.W.
WASHINGTON, DC 20005-3315

202 • 408 • 4000
FACSIMILE 202 • 408 • 4400

ATLANTA
404•653•6400
PALO ALTO
650•849•6600

WRITER'S DIRECT DIAL NUMBER:

(202) 408-4470

A
TOKYO
011•813•3431•6943
BRUSSELS
011•322•646•0353

JC678 U.S. PTO
09/457889
12/10/99

December 10, 1999

Assistant Commissioner for Patents
Washington, DC 20231

Re: New U.S. Patent Application
Title: METHOD AND SYSTEM FOR FACILITATING
RELOCATION OF DEVICES ON A NETWORK
Inventors: Germano Caronni, Amit Gupta, Sandeep Kumar,
Tom R. Markson, Christoph L. Schuba, and Glenn C. Scott
ATTORNEY DOCKET NO. 6502.0283

Sir:

We enclose the following papers for filing in the United States Patent and Trademark Office in connection with the above patent application.

1. Application - 37 pages, including 6 independent claims and 24 claims total.
2. Drawings - 12 sheets of informal drawings (Figs. 1 - 11).

This application is being filed under the provisions of 37 C.F.R. §§ 1.53(b) and 1.53(f). Applicants await notification from the Patent and Trademark Office of the time set for filing the Declaration.

Please accord this application a serial number and filing date.

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER, L.L.P.

Assistant Commissioner of Patents
December 10, 1999
Page 2

No filing fee is being paid at this time. With the exception of the filing fee, the Commissioner is hereby authorized to charge any fees due, including fees under 37 C.F.R. § 1.16 or § 1.17, during the pendency of this application to our Deposit Account No. 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER

By


Michael L. Kiklis
Reg. No. 38,939

MLK:laf
Enclosures

UNITED STATES PATENT APPLICATION

FOR

METHOD AND SYSTEM FOR FACILITATING RELOCATION
OF DEVICES ON A NETWORK

BY

GERMANO CARONNI

AMIT GUPTA

SANDEEP KUMAR

TOM R. MARKSON

CHRISTOPH L. SCHUBA

GLENN C. SCOTT

RELATED APPLICATIONS

The following identified U.S. patent applications are relied upon and are incorporated by reference in this application.

U.S. Patent Application No. _____, entitled "SYSTEM AND METHOD FOR SEPARATING ADDRESSES FROM THE DELIVERY SCHEME IN A VIRTUAL PRIVATE NETWORK," bearing attorney docket no. 06502.0280, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "TRULY ANONYMOUS COMMUNICATIONS USING SUPERNETS WITH THE PROVISION OF TOPOLOGY HIDING," bearing attorney docket no. 06502.0281, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "SANDBOXING APPLICATIONS IN A PRIVATE NETWORK USING A PUBLIC-NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0284, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "SECURE ADDRESS RESOLUTION FOR A PRIVATE NETWORK USING A PUBLIC NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0285, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "DECOUPLING ACCESS CONTROL FROM KEY MANAGEMENT IN A NETWORK," bearing attorney docket no. 06502.0286, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "CHANNEL-SPECIFIC FILE SYSTEM VIEWS IN A PRIVATE NETWORK USING A PUBLIC NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0287, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "PRIVATE NETWORK USING A PUBLIC-NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0288, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "SYSTEM AND METHOD FOR ENABLING SCALABLE SECURITY IN A VIRTUAL PRIVATE NETWORK," bearing attorney docket no. 06502.0289, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "USING MULTICASTING TO PROVIDE ETHERNET-LIKE COMMUNICATION BEHAVIOR TO SELECTED PEERS ON A NETWORK," bearing attorney docket no. 06502.0290, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "ANYCASTING IN A PRIVATE NETWORK USING A PUBLIC NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0291, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "SCALABLE SECURITY ASSOCIATIONS FOR GROUPS FOR USE IN A PRIVATE NETWORK USING A PUBLIC-NETWORK INFRASTRUCTURE," bearing attorney docket no. 06502.0292, and filed on the same date herewith.

U.S. Patent Application No. _____, entitled "ENABLING SIMULTANEOUS PROVISION OF INFRASTRUCTURE SERVICES," bearing attorney docket no. 06502.0293, and filed on the same date herewith.

DRAFT - DO NOT CITE

FIELD OF THE INVENTION

The present invention relates generally to data processing systems and, more particularly, to a private network using a public-network infrastructure.

BACKGROUND OF THE INVENTION

As part of their day-to-day business, many organizations require an enterprise network, a private network with lease lines, dedicated channels, and network connectivity devices, such as routers, switches, and bridges. These components, collectively known as the network's "infrastructure," are very expensive and require a staff of information technology personnel to maintain them. This maintenance requirement is burdensome on many organizations whose main business is not related to the data processing industry (e.g., a clothing manufacturer) because they are not well suited to handle such data processing needs.

Another drawback to enterprise networks is that they are geographically restrictive. The term "geographically restrictive" refers to the requirement that if a user is not physically located such that they can plug their device directly into the enterprise network, the user cannot typically utilize it. To alleviate the problem of geographic restrictiveness, virtual private networks have been developed.

In a virtual private network (VPN), a remote device or network connected to the Internet may connect to the enterprise network through a firewall. This allows the remote device to access resources on the enterprise network even though it may not be located near any component of the enterprise network. For example, Fig. 1 depicts a VPN 100, where

enterprise network 102 is connected to the Internet 104 via firewall 106. By using VPN 100, a remote device D₁ 108 may communicate with enterprise network 102 via Internet 104 and firewall 106. Thus, D₁ 108 may be plugged into an Internet portal virtually anywhere within the world and make use of the resources on enterprise network 102.

5 To perform this functionality, D₁ 108 utilizes a technique known as tunneling to ensure that the communication between itself and enterprise network 102 is secure in that it cannot be viewed by an interloper. "Tunneling" refers to encapsulating one packet inside another when packets are transferred between end points (e.g., D₁ 108 and VPN software 109 running on firewall 106). The packets may be encrypted at their origin and decrypted at their 10 destination. For example, Fig. 2A depicts a packet 200 with a source Internet protocol (IP) address 202, a destination IP address 204, and data 206. It should be appreciated that packet 200 contains other information not depicted, such as the source and destination port. As shown in Fig. 2B, the tunneling technique forms a new packet 208 out of packet 200 by encrypting it and adding both a new source IP address 210 and a new destination IP address 15 212. In this manner, the contents of the original packet (i.e., 202, 204, and 206) are not visible to any entity other than the destination. Referring back to Fig. 1, by using tunneling, remote device D₁ 108 may communicate and utilize the resources of the enterprise network 102 in a secure manner.

Although VPNs alleviate the problem of geographic restrictiveness, they impose 20 significant processing overhead when two remote devices communicate. For example, if remote device D₁ 108 wants to communicate with remote device D₂ 110, D₁ sends a packet using tunneling to VPN software 109, where the packet is decrypted and then transferred to

the enterprise network 102. Then, the enterprise network 102 sends the packet to VPN software 109, where it is encrypted again and transferred to D₂. Given this processing overhead, it is burdensome for two remote devices to communicate in a VPN environment. It is therefore desirable to alleviate the need of organizations to maintain their own network infrastructure as well as to improve communication between remote devices.

5

SUMMARY OF THE INVENTION

Methods and systems consistent with the present invention provide a private network that uses components from a public-network infrastructure. Nodes of the private network can be located on virtually any device in the public network (e.g., the Internet), and both their communication and utilization of resources occur in a secure manner. As a result, the users of this private network benefit from their network infrastructure being maintained for them as part of the public-network infrastructure, while the level of security they receive is similar to or even stronger than that provided by conventional private networks. Additionally, the nodes of the private network are not geographically restricted in that they can be connected to the private network from virtually any portal to the Internet in the world.

This private network also provides flexible and dynamic mobility support. Sometimes, the device on which a node runs is relocated to a new physical location (e.g., a new office). In this situation, a problem arises because the nodes that send communications to the moving node will be unable to do so once the moving node relocates. This problem occurs because when the device moves, nodes that run on that device receive a new IP address. Some conventional systems solve this problem by using a proxy as a middleman

10

15

20

between the source node and the destination node. In these systems, the source node sends a packet to the proxy, and the proxy then sends it to the destination node. Then, when the destination node moves, it updates the proxy with its new address so that it can continue to receive communications. Such systems incur significant processing overhead because of use
5 of the proxy. The private network according to an implementation of the present invention does not use a proxy; instead, the private network sends communications directly from the sending node to the destination node. When the destination node moves, the destination node updates its address with the sending nodes that communicate with it so that point-to-point communication can resume.

10 In one implementation of the present invention, a method is provided in a data processing system for providing point-to-point communication in a network with a source node and a destination node. The source node accesses the address of the destination node and sends a first packet to the destination node using the address. The destination node receives the first packet. The address of the destination node is updated to a new address in
15 response to a change in the destination node's address to the new address. The source node sends a second packet to the destination node using the new address, and the destination node receives the second packet.

In accordance with an implementation consistent with the present invention, a plurality of devices are connected to a distributed system. A first of the devices comprises
20 a memory and a processor, the memory having a source node that sends a first packet to a destination node using an address of the destination node, receives a new address to supersede the address of the destination node, responsive to a change in the address of the

destination node to the new address, and sends a second packet to the destination node at the new address. The second device comprises a memory and a processor, the memory having the destination node, that receives the first packet at the address, sends the new address to the source node in response to the change in the address of the destination node to the new address, and receives the second packet.

5

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

10

Fig. 1 depicts a conventional virtual private network (VPN) system;

Fig. 2A depicts a conventional network packet;

Fig. 2B depicts the packet of Fig. 2A after it has been encrypted in accordance with a conventional tunneling technique;

15

Fig. 3 depicts a data processing system suitable for use with methods and systems consistent with the present invention;

Fig. 4 depicts the nodes depicted in Fig. 3 communicating over multiple channels;

Fig. 5 depicts two devices depicted in Fig. 3 in greater detail;

Figs. 6A and 6B depict a flow chart of the steps performed when a node joins a VPN in a manner consistent with the present invention;

20

Fig. 7 depicts a flow chart of the steps performed when sending a packet from a node of the VPN in a manner consistent with the present invention;

Fig. 8 depicts a flow chart of the steps performed when receiving a packet by a node of the VPN in a manner consistent with the present invention;

Fig. 9 depicts a flow chart of the steps performed when logging out of a VPN in a manner consistent with the present invention;

5 Fig. 10 depicts a flow chart of the steps performed to maintain point-to-point communication when a node's IP address changes; and

Fig. 11 depicts a flow chart of the steps performed when updating a destination node's address mapping with nodes that communicate with it.

DETAILED DESCRIPTION

10 Methods and systems consistent with the present invention provide a "Supernet," which is a private network that uses components from a public-network infrastructure. A Supernet allows an organization to utilize a public-network infrastructure for its enterprise network so that the organization no longer has to maintain a private network infrastructure; instead, the organization may have the infrastructure maintained for them by one or more service providers or other organizations that specialize in such connectivity matters. As such, the burden of maintaining an enterprise network is greatly reduced. Moreover, a Supernet is not geographically restrictive, so a user may plug their device into the Internet from virtually any portal in the world and still be able to use the resources of their private network in a secure and robust manner.

15 20 The Supernet provides flexible and dynamic mobility support, because when a destination node moves to a new location and receives a new IP address, the destination node

automatically updates the sending nodes with its new IP address. Thus, a node can change locations repeatedly and continue to communicate directly with other nodes without the use of a proxy or other middleman as is used in some conventional systems.

Overview

Fig. 3 depicts a data processing system 300 suitable for use with methods and systems consistent with the present invention. Data processing system 300 comprises a number of devices, such as computers 302-312, connected to a public network, such as the Internet 314. A Supernet's infrastructure uses components from the Internet because devices 302, 304, and 312 contain nodes that together form a Supernet and that communicate by using the infrastructure of the Internet. These nodes 316, 318, 320, and 322 are communicative entities (e.g., processes) running within a particular device and are able to communicate among themselves as well as access the resources of the Supernet in a secure manner. When communicating among themselves, the nodes 316, 318, 320, and 322 serve as end points for the communications, and no other processes or devices that are not part of the Supernet are able to communicate with the Supernet's nodes or utilize the Supernet's resources. The Supernet also includes an administrative node 306 to administer to the needs of the Supernet.

It should be noted that since the nodes of the Supernet rely on the Internet for connectivity, if the device on which a node is running relocates to another geographic location, the device can be plugged into an Internet portal and the node running on that device can quickly resume the use of the resources of the Supernet. It should also be noted that since a Supernet is layered on top of an existing network, it operates independently of the transport layer. Thus, the nodes of a Supernet may communicate over different

transports, such as IP, IPX, X.25, or ATM, as well as different physical layers, such as RF communication, cellular communication, satellite links, or land-based links.

As shown in Fig. 4, a Supernet includes a number of channels that its nodes 316-322 can communicate over. A "channel" refers to a collection of virtual links through the public-network infrastructure that connect the nodes on the channel such that only these nodes can communicate over it. A node on a channel may send a message to another node on that channel, known as a unicast message, or it can send a message to all other nodes on that channel, known as a multicast message. For example, channel 1 402 connects node A 316 and node C 320, and channel 2 404 connects node B 318, node C 320, and node D 322. Each Supernet has any number of preconfigured channels over which the nodes on that channel can communicate. In an alternative embodiment, the channels are dynamically defined.

In addition to communication, the channels may be used to share resources. For example, channel 1 402 may be configured to share a file system as part of node C 320 such that node A 316 can utilize the file system of node C in a secure manner. In this case, node C 320 serves as a file system manager by receiving file system requests (e.g., open, close, read, write, etc.) and by satisfying the requests by manipulating a portion of the secondary storage on its local machine. To maintain security, node C 320 stores the data in an encrypted form so that it is unreadable by others. Such security is important because the secondary storage may not be under the control of the owners of the Supernet, but may instead be leased from a service provider. Additionally, channel 2 404 may be configured to share the computing resources of node D 322 such that nodes B 318 and C 320 send code

to node D for execution. By using channels in this manner, resources on a public network can be shared in a secure manner.

A Supernet provides a number of features to ensure secure and robust communication among its nodes. First, the system provides authentication and admission control so that 5 nodes become members of the Supernet under strict control to prevent unauthorized access.

Second, the Supernet provides communication security services so that the sender of a message is authenticated and communication between end points occurs in a secure manner by using encryption. Third, the system provides key management to reduce the possibility of an intruder obtaining an encryption key and penetrating a secure communication session.

10 The system does so by providing one key per channel and by changing the key for a channel whenever a node joins or leaves the channel. Alternatively, the system may use a different security policy.

Fourth, the system provides address translation in a transparent manner. Since the Supernet is a private network constructed from the infrastructure of another network, the 15 Supernet has its own internal addressing scheme, separate from the addressing scheme of the underlying public network. Thus, when a packet from a Supernet node is sent to another Supernet node, it travels through the public network. To do so, the Supernet performs address translation from the internal addressing scheme to the public addressing scheme and vice versa. To reduce the complexity of Supernet nodes, system-level components of the 20 Supernet perform this translation on behalf of the individual nodes so that it is transparent to the nodes. Another benefit of the Supernet's addressing is that it uses an IP-based internal

addressing scheme so that preexisting programs require little modification to run within a Supernet.

Fifth, the Supernet provides operating system-level enforcement of node compartmentalization in that an operating system-level component treats a Supernet node running on a device differently than it treats other processes on that device. This component (i.e., a security layer in a protocol stack) recognizes that a Supernet node is part of a Supernet, and therefore, it enforces that all communications to and from this node travel through the security infrastructure of the Supernet such that this node can communicate with other members of the Supernet and that non-members of the Supernet cannot access this node. Additionally, this operating system-level enforcement of node compartmentalization allows more than one Supernet node to run on the same machine, regardless of whether the nodes are from the same Supernet, and allows nodes of other networks to run on the same machine as a Supernet node.

Finally, the Supernet provides mobility support to its nodes in a flexible and dynamic fashion. To send a packet form a source node to a destination node in a conventional network, the packet typically travels through a proxy acting as a middleman. The source node sends the packet to the proxy using the proxy's address, and upon receipt, the proxy inserts the destination node's IP address into the packet and sends it to the destination node.. Supernets eliminate the need for the proxy. In a Supernet, the source node and the destination node communicate in a point-to-point manner, and when the destination node moves to a new location, it automatically updates the sending nodes with its new IP address, thus maintaining the point-to-point communication. The destination node can choose from

a number of ways to update the sending nodes. As a result, a Supernet provides flexible and dynamic mobility support, allowing a node to change locations repeatedly and continue to communicate directly with other nodes.

Implementation Details

5 Fig. 5 depicts administrative machine 306 and device 302 in greater detail, although the other devices 304 and 308-312 may contain similar components. Device 302 and administrative machine 306 communicate via Internet 314. Each device contains similar components, including a memory 502, 504; secondary storage 506, 508; a central processing unit (CPU) 510, 512; an input device 514, 516; and a video display 518, 520. One skilled in the art will appreciate that these devices may contain additional or different components.
10

Memory 504 of administrative machine 306 includes the SASD process 540, VARPDB 548, and KMS 550 all running in user mode. That is, CPU 512 is capable of running in at least two modes: user mode and kernel mode. When CPU 512 executes programs running in user mode, it prevents them from directly manipulating the hardware components, such 15 as video display 518. On the other hand, when CPU 512 executes programs running in kernel mode, it allows them to manipulate the hardware components. Memory 504 also contains a VARPDB 551 and a TCP/IP protocol stack 552 that are executed by CPU 512 running in kernel mode. TCP/IP protocol stack 552 contains a TCP/UDP layer 554 and an IP layer 556, both of which are standard layers well known to those of ordinary skill in the 20 art. Secondary storage 508 contains a configuration file 558 that stores various configuration-related information (described below) for use by SASD 540.

SASD 540 represents a Supernet: there is one instance of an SASD per Supernet, and it both authenticates nodes and authorizes nodes to join the Supernet. VARPД 548 has an associated component, VARPDB 551, into which it stores mappings of the internal Supernet addresses, known as a node IDs, to the network addresses recognized by the public-network infrastructure, known as the real addresses. The "node ID" may include the following: a Supernet ID (e.g., 0x123), reflecting a unique identifier of the Supernet, and a virtual address, comprising an IP address (e.g., 10.0.0.1). The "real address" is an IP address (e.g., 10.0.0.2) that is globally unique and meaningful to the public-network infrastructure. In a Supernet, one VARPД runs on each machine, and it may play two roles. First, a VARPД may act as a server by storing all address mappings for a particular Supernet into its associated VARPDB. Second, regardless of its role as a server or not, each VARPД assists in address translation for the nodes on its machine. In this role, the VARPД stores into its associated VARPDB the address mappings for its nodes, and if it needs a mapping that it does not have, it will contact the VARPД that acts as the server for the given Supernet to obtain it.

KMS 550 performs key management by generating a new key every time a node joins a channel and by generating a new key every time a node leaves a channel. There is one KMS per channel in a Supernet.

To configure a Supernet, a system administrator creates a configuration file 558 that is used by SASD 540 when starting or reconfiguring a Supernet. This file may specify: (1) the Supernet name, (2) all of the channels in the Supernet, (3) the nodes that communicate over each channel, (4) the address of the KMS for each channel, (5) the address of the VARPД that acts as the server for the Supernet, (6) the user IDs of the users who are

authorized to create Supernet nodes, (7) the authentication mechanism to use for each user of each channel, and (8) the encryption algorithm to use for each channel. Although the configuration information is described as being stored in a configuration file, one skilled in the art will appreciate that this information may be retrieved from other sources, such as
5 databases or interactive configurations.

After the configuration file is created, it is used to start a Supernet. For example, when starting a Supernet, the system administrator first starts SASD, which reads the configuration information stored in the configuration file. Then, the administrator starts the VARPD on the administrator's machine, indicating that it will act as the server for the
10 Supernet and also starts the KMS process. After this processing has completed, the Supernet is ready for nodes to join it.

Memory 502 of device 302 contains SNlogin script 522, SNlogout script 524, VARPD 526, KMC 528, KMD 530, and node A 522, all running in user mode. Memory 502 also includes TCP/IP protocol stack 534 and VARPDB 536 running in kernel mode.

15 SNlogin 522 is a script used for logging into a Supernet. Successfully executing this script results in a Unix shell from which programs (e.g., node A 522) can be started to run within the Supernet context, such that address translation and security encapsulation is performed transparently for them and all they can typically access is other nodes on the Supernet. Alternatively, a parameter may be passed into SNlogin 522 that indicates a
20 particular process to be automatically run in a Supernet context. Once a program is running in a Supernet context, all programs spawned by that program also run in the Supernet context, unless explicitly stated otherwise. SNlogout 524 is a script used for logging out of

a Supernet. Although both SNlogin 522 and SNlogout 524 are described as being scripts, one skilled in the art will appreciate that their processing may be performed by another form of software. VARPD 526 performs address translation between node IDs and real addresses. KMC 528 is the key management component for each node that receives updates whenever
5 the key for a channel ("the channel key") changes. There is one KMC per node per channel. KMD 530 receives requests from SNSL 542 of the TCP/IP protocol stack 534 when a packet is received and accesses the appropriate KMC for the destination node to retrieve the appropriate key to decrypt the packet. Node A 532 is a Supernet node running in a Supernet context.

10 TCP/IP protocol stack 534 contains a standard TCP/UDP layer 538, two standard IP layers (an inner IP layer 540 and an outer IP layer 544), and a Supernet security layer (SNSL) 542, acting as the conduit for all Supernet communications. To conserve memory, both inner IP layer 540 and outer IP layer 544 may share the same instance of the code of an IP layer. SNSL 542 performs security functionality as well as address translation. It also caches the
15 most recently used channel keys for ten seconds. Thus, when a channel key is needed, SNSL 542 checks its cache first, and if it is not found, it requests KMD 530 to contact the appropriate KMC to retrieve the appropriate channel key. Two IP layers 540, 544 are used in the TCP/IP protocol stack 534 because both the internal addressing scheme and the external addressing scheme are IP-based. Thus, for example, when a packet is sent, inner IP
20 layer 540 receives the packet from TCP/UDP layer 538 and processes the packet with its node ID address before passing it to the SNSL layer 542, which encrypts it, prepends the real

source IP address and the real destination IP address, and then passes the encrypted packet to outer IP layer 544 for sending to the destination.

SNSL 542 utilizes VARPDB 536 to perform address translation. VARPDB stores all of the address mappings encountered thus far by SNSL 542. If SNSL 542 requests a mapping that VARPDB 536 does not have, VARPDB communicates with the VARPD 526 on the local machine to obtain the mapping. VARPD 526 will then contact the VARPD that acts as the server for this particular Supernet to obtain it.

Although aspects of the present invention are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on or read from other types of computer-readable media, such as secondary storage devices, like hard disks, floppy disks, or CD-ROM; a carrier wave from a network, such as the Internet; or other forms of RAM or ROM either currently known or later developed. Additionally, although a number of the software components are described as being located on the same machine, one skilled in the art will appreciate that these components may be distributed over a number of machines.

Figs. 6A and 6B depict a flow chart of the steps performed when a node joins a Supernet. The first step performed is that the user invokes the SNlogin script and enters the Supernet name, their user ID, their password, and a requested virtual address (step 602). Of course, this information depends on the particular authentication mechanism used. Upon receiving this information, the SNlogin script performs a handshaking with SASD to authenticate this information. In this step, the user may request a particular virtual address to be used, or alternatively, the SASD may select one for them. Next, if any of the

information in step 602 is not validated by SASD (step 604), processing ends. Otherwise, upon successful authentication, SASD creates an address mapping between a node ID and the real address (step 606). In this step, SASD concatenates the Supernet ID with the virtual address to create the node ID, obtains the real address of the SNlogin script by querying network services in a well-known manner, and then registers this information with the VARPD that acts as the server for this Supernet. This VARPD is identified in the configuration file.

After creating the address mapping, SASD informs the KMS that there is a new Supernet member that has been authenticated and admitted (step 608). In this step, SASD sends the node ID and the real address to KMS who then generates a key ID, a key for use in communicating between the node's KMC and the KMS ("a node key"), and updates the channel key for use in encrypting traffic on this particular channel (step 610). Additionally, KMS sends the key ID and the node key to SASD and distributes the channel key to all KMCs on the channel as a new key because a node has just been added to the channel.

SASD receives the key ID and the node key from KMS and returns it to SNlogin (step 612). After receiving the key ID and the node key from SASD, SNlogin starts a KMC for this node and transmits to the KMC the node ID, the key ID, the node key, the address of the VARPD that acts as the server for this Supernet, and the address of KMS (step 614). The KMC then registers with the KMD indicating the node it is associated with, and KMC registers with KMS for key updates (step 616). When registering with KMS, KMC provides its address so that it can receive updates to the channel key via the Versakey protocol. The Versakey protocol is described in greater detail in IEEE Journal on Selected Areas in Communication,

S E C U R I T Y
C O M P L I A N C E
A U T O M A T I O N
P R O T E C T I O N

Vol. 17, No. 9, 1999, pp. 1614-1631. After registration, the KMC will receive key updates whenever a channel key changes on one of the channels that the node communicates over.

Next, SNlogin configures SNSL (step 618 in Fig. 6B). In this step, SNlogin indicates which encryption algorithm to use for this channel and which authentication algorithm to use, both of which are received from the configuration file via SASD. SNSL stores this information in an access control list. In accordance with methods and systems consistent with present invention, any of a number of well-known encryption algorithms may be used, including the Data Encryption Standard (DES), Triple-DES, the International Data Encryption Algorithm (IDEA), and the Advanced Encryption Standard (AES). Also, RC2, RC4, and RC5 from RSA Incorporated may be used as well as Blowfish from Counterpane.com. Additionally, in accordance with methods and systems consistent with the present invention, any of a number of well-known authentication algorithms may be used, including Digital Signatures, Kerberos, Secure Socket Layer (SSL), and MD5, which is described in RFC1321 of the Internet Engineering Task Force, April, 1992.

After configuring SNSL, SNlogin invokes an operating system call, SETVIN, to cause the SNlogin script to run in a Supernet context (step 620). In Unix, each process has a data structure known as the "proc structure" that contains the process ID as well as a pointer to a virtual memory description of this process. In accordance with methods and systems consistent with the present invention, the channel IDs indicating the channels over which the process communicates as well as its virtual address for this process are added to this structure. By associating this information with the process, the SNSL layer can enforce that this process runs in a Supernet context. Although methods and systems consistent with the

present invention are described as operating in a Unix environment, one skilled in the art will appreciate that such methods and systems can operate in other environments. After the SNlogin script runs in the Supernet context, the SNlogin script spawns a Unix program, such as a Unix shell or a service daemon (step 622). In this step, the SNlogin script spawns a
5 Unix shell from which programs can be run by the user. All of these programs will thus run in the Supernet context until the user runs the SNlogout script.

Fig. 7 depicts a flow chart of the steps performed when sending a packet from node A. Although the steps of the flow chart are described in a particular order, one skilled in the art will appreciate that these steps may be performed in a different order. Additionally,
10 although the SNSL layer is described as performing both authentication and encryption, this processing is policy driven such that either authentication, encryption, both, or neither may be performed. The first step performed is for the SNSL layer to receive a packet originating from node A via the TCP/UDP layer and the inner IP layer (step 702). The packet contains a source node ID, a destination node ID, and data. The SNSL layer then accesses the
15 VARPDB to obtain the address mapping between the source node ID and the source real address as well as the destination node ID and the destination real address (step 704). If they are not contained in the VARPDB because this is the first time a packet has been sent from this node or sent to this destination, the VARPDB accesses the local VARPD to obtain the mapping. When contacted, the VARPD on the local machine contacts the VARPD that acts
20 as the server for the Supernet to obtain the appropriate address mapping.

After obtaining the address mapping, the SNSL layer determines whether it has been configured to communicate over the appropriate channel for this packet (step 706). This

configuration occurs when SNlogin runs, and if the SNSL has not been so configured, processing ends. Otherwise, SNSL obtains the channel key to be used for this channel (step 708). The SNSL maintains a local cache of keys and an indication of the channel to which each key is associated. Each channel key is time stamped to expire in ten seconds, although
5 this time is configurable by the administrator. If there is a key located in the cache for this channel, SNSL obtains the key. Otherwise, SNSL accesses KMD which then locates the appropriate channel key from the appropriate KMC. After obtaining the key, the SNSL layer encrypts the packet using the appropriate encryption algorithm and the key previously obtained (step 710). When encrypting the packet, the source node ID, the destination node
10 ID, and the data may be encrypted, but the source and destination real addresses are not, so that the real addresses can be used by the public network infrastructure to send the packet to its destination.

After encrypting the packet, the SNSL layer authenticates the sender to verify that it is the bona fide sender and that the packet was not modified in transit (step 712). In this step,
15 the SNSL layer uses the MD5 authentication protocol, although one skilled in the art will appreciate that other authentication protocols may be used. Next, the SNSL layer passes the packet to the IP layer where it is then sent to the destination node in accordance with known techniques associated with the IP protocol (step 714).

Fig. 8 depicts a flow chart of the steps performed by the SNSL layer when it receives
20 a packet. Although the steps of the flow chart are described in a particular order, one skilled in the art will appreciate that these steps may be performed in a different order. Additionally, although the SNSL layer is described as performing both authentication and encryption, this

processing is policy driven such that either authentication, encryption, both, or neither may be performed. The first step performed by the SNSL layer is to receive a packet from the network (step 801). This packet contains a real source address and a real destination address that are not encrypted as well as a source node ID, a destination node ID, and data that are encrypted. Then, it determines whether it has been configured to communicate on this channel to the destination node (step 802). If SNSL has not been so configured, processing ends. Otherwise, the SNSL layer obtains the appropriate key as previously described (step 804). It then decrypts the packet using this key and the appropriate encryption algorithm (step 806). After decrypting the packet, the SNSL layer authenticates the sender and validates the integrity of the packet (step 808), and then it passes the packet to the inner IP layer for delivery to the appropriate node (step 810). Upon receiving the packet, the inner IP layer uses the destination node ID to deliver the packet.

Fig. 9 depicts a flow chart of the steps performed when logging a node out of a Supernet. The first step performed is for the user to run the SNlogout script and to enter a node ID (step 902). Next, the SNlogout script requests a log out from SASD (step 904). Upon receiving this request, SASD removes the mapping for this node from the VARPDL that acts as the server for the Supernet (step 906). SASD then informs KMS to cancel the registration of the node, and KMS terminates this KMC (step 908). Lastly, KMS generates a new channel key for the channels on which the node was communicating (step 910) to provide greater security.

Mobility Support

Fig. 10 depicts a flow chart of the steps performed to maintain point-to-point communication when a node's IP address changes. This process occurs when the device on which a node operates changes location (step 1000). After it moves, the device contacts the Internet Assigned Numbers Authority (IANA) to receive a new IP address (step 1002). Next, each destination node operating on the device updates its address mapping with its new IP address (step 1004). This step can be performed a number of ways as is further described below. Next, a sending node sends a packet to the destination node using the new IP address (step 1008), and the destination node receives the packet at its new location (step 1008).

Fig. 11 depicts a flow chart of one way a destination node can update the sending nodes after it moves. First, the destination node sends its new IP address to the VARPD acting as the server for each channel that it uses (step 1102). Each VARPD server updates its VARPDB by replacing the destination node's old IP address with its new IP address in the address mapping (step 1104). The VARPD server then sends a message to the local VARPD for each node on the channel (the sending nodes) with the destination node's new mapping (step 1106). Each local VARPD receives the message and updates its own VARPDB by replacing the destination node's old IP address with its new IP address in the mapping (step 1108).

Alternatively, the destination node performs the updating step by sending a message containing its new IP address to each node with an address mapping stored in the destination node's VARPDB. The nodes with mappings stored in the destination node's VARPDB are those with whom the destination node has recently communicated. Each sending node

receives the message and updates its own VARPDB by replacing the destination node's old IP address with its new IP address in the mapping. The destination node then sends its new IP address to the VARPD acting as the server for each channel it uses. Each server updates its VARPDB by replacing the destination node's old IP address with its new IP address in the address mapping.

In another alternative, the destination node can use multicasting to perform the updating step in either a proactive or reactive manner. Nodes in the Supernet implement multicasting using the well-known Internet Group Management Protocol (IGMP). To be proactive, the destination node sends a message with its new IP address to the multicast address for each channel it uses, thus notifying each node on the channel. The Supernet uses existing Internet routing tables to send the message to each node on the channel. Each sending node receives the message and updates its own VARPDB by replacing the destination node's old IP address with its new IP address in the mapping.

To be reactive, the destination node sends a message with its new IP address to the closest Internet router requesting that it be added to the multicast group for each channel it uses and sends a unicast message to the VARPD server providing it with the new address. The destination node then waits until a sending node sends a message to the multicast address for the channel it shares with the destination node. The destination node receives the message and sends a unicast message containing the destination node's new IP address to the sending node. The sending node receives the message and updates its own VARPDB by replacing the destination node's old IP address with its new IP address in the mapping. The destination node does this until it sends its new address to each node on the channel.

Although the present invention has been described with reference to a preferred embodiment, those skilled in the art will know of various changes in form and detail which may be made without departing from the spirit and scope of the present invention as defined in the appended claims and their full scope of equivalents.

U.S. PATENT AND TRADEMARK OFFICE
SEARCHED
INDEXED
MAILED
FILED
JULY 19 1999

LAW OFFICES
FINNEMAN, HENDERSON,
FARABOW, GARRETT,
& DUNNER, L.L.P.
1300 I STREET, N.W.
WASHINGTON, D.C. 20005
202-408-4000

CLAIMS

What is claimed is:

1. A method in a distributed system for communicating in a network with a source node and a destination node, wherein the source node and the destination node have an address, the method comprising the steps of:

accessing the address of the destination node by the source node, wherein the
5 source node and the destination node are programs;

sending a first packet by the source node to the destination node by using the accessed address;

receiving the first packet by the destination node;

10 updating the accessed address to reflect a new address of the destination node responsive to a change in the address of the destination node to the new address;

sending a second packet by the source node to the destination node by using the new address; and

receiving the second packet by the destination node.

15 2. The method of claim 1, wherein the source node and the destination node have a local address cache, and wherein the receiving the first packet step includes the steps of:

storing, in the local cache of the destination node, the address of the sending node;
and

wherein the updating step further includes the steps of:

RECEIVED
U.S. PATENT AND TRADEMARK OFFICE
JULY 19 1993

retrieving from the local cache of the destination node the address of the sending node;

sending a third packet containing the new address of the destination node to the source node by the destination node using the address of the source node;

5 receiving the third packet by the source node; and

storing, in the local cache of the source node, the new address of the destination node.

3. The method of claim 1, wherein the distributed system has a central address store,

10 the method further comprising the steps of:

storing the address of the source node and the destination node in the central address store; and

wherein the sending a first packet step further includes the step of:

accessing the address of the destination node from the central address 15 store; and

wherein the updating step further includes the steps of:

sending a third packet containing the new address of the destination node to the central address store by the destination node; and

storing the new address of the destination node in the central address store;

20 and

accessing the new address by the source node.

4. The method of claim 1, wherein the source node and the destination node communicate with other nodes by using a multicast address such that a communication sent to the multicast address is sent to the source node, the destination node, and the other nodes, the source node and the other nodes having a local cache, wherein the updating step further includes the steps of:

sending a third packet containing the new address of the destination node to the source node by the destination node using the multicast address;

receiving the third packet by the source node and the other nodes; and

storing in the local cache of the source node and the other nodes the new address of the destination node.

5. The method of claim 1, wherein the source node and the destination node have a local address cache and communicate by using a multicast address, and wherein the updating step further includes the steps of:

sending a joining request, by the destination node, to a router to add the new address of the destination node to the multicast group;

sending a message, by the source node, to the destination node via the multicast address;

receiving the message by the destination node;

sending a third packet containing the new address of the destination node by the destination node to the source node, using the address for the source node;

receiving the third packet by the source node; and

storing the new address of the destination node in the local cache of the source node.

6. The method of claim 1 wherein the network is a private network running on a public network infrastructure.

5 7. The method of claim 1 wherein the change in the address of the destination node to the new address is caused by a device on which the destination node runs physically changing locations.

8. A method in a distributed system for communicating in a network with a source node and a destination node, wherein the source node and the destination node have an address, the method comprising the steps of:

receiving a first packet by the destination node from the source node, the packet being addressed to the address of the destination node, wherein the source node and the destination node are programs;

15 updating the address of the destination node to a new address responsive to a change in the address of the destination node to the new address; and

receiving a second packet by the destination node at the new address.

9. The method of claim 8, wherein the destination node has a local address cache, and wherein the receiving first packet step includes the steps of:

storing, in the local cache of the destination node, the address of the sending node;
and

wherein the updating step further includes the steps of:

5 retrieving from the local cache of the destination node the address of the
sending node; and

sending a third packet containing the new address of the destination node
by the destination node to the source node, using the address of the source node.

10. The method of claim 8, wherein the distributed system has a central address store,
10 the method further comprising the steps of:

storing the address of the source node and the destination node in the central
address store; and

wherein the sending a first packet step further includes the step of:

accessing the address of the destination node from the central address
15 store; and

wherein the updating step further includes the step of:

sending a third packet containing the new address of the destination node
to the central address store by the destination node .

11. The method of claim 8, wherein the source node and the destination node
20 communicate with other nodes by using a multicast address such that a communication
sent to the multicast address is sent to the source node, the destination node, and the other

nodes, the source node and the other nodes having a local cache, and wherein the updating step further includes the step of:

sending a third packet containing the new address of the destination node to the source node by the destination node using the multicast address.

5 12. The method of claim 8, wherein the source node and the destination node have a local address cache and communicate by using a multicast address, and wherein the updating step further includes the steps of:

sending a joining request, by the destination node, to a router to add the new address of the destination node to the multicast group;

10 receiving a message from the source node by the destination node, via the multicast address; and

sending a third packet containing the new address of the destination node by the destination node to the source node, using the address for the source node.

15 13. The method of claim 8 wherein the network is a private network running on a public network infrastructure.

14. The method of claim 8 wherein the change in the address of the destination node to the new address is caused by a device on which the destination node runs physically changing locations.

15. A method in a data processing system for communicating in a network with a source node and a destination node, wherein each node has an address, the method comprising the steps of:

sending a first packet from the source node to the destination node by using the
5 address of the destination node, wherein the source node and the destination node are programs;

receiving a new address by the source node to supersede the address of the destination node responsive to a change in the address of the destination node to the new
10 address; and

10 sending a second packet from the source node to the destination node by using the new address.

16. The method of claim 15, wherein the network is a private network running on a public network infrastructure.

17. The method of claim 15 wherein the change in the address of the destination node
15 to the new address is caused by a device on which the destination node runs physically changing locations.

18. A distributed system with a plurality of devices, comprising:
a first of the devices comprising:

a memory with a source node that sends a first packet to a destination node using an address of the destination node, that receives a new address to supersede the address of the destination node responsive to a change in the address of the destination node to the new address, and that sends a second packet to the destination node using the
5 new address;

a processor that runs the source node; and

a second device comprising:

a memory with the destination node that receives the first packet at the address and that sends the new address to the source node in response to the change in the
10 address of the destination node to the new address, and that receives the second packet;
and

a processor for running the destination node.

19. The distributed system of claim 18 wherein the network is a private network running on a public network infrastructure.

15 20. The distributed system of claim 18 wherein the change in the address of the destination node to the new address is caused by a device on which the destination node runs physically changing locations.

21. A data processing system for communicating in a network with a source node and a destination node, wherein the source node and the destination node have an address, the data processing system comprising:

means for accessing the address of the destination node;

means for sending a first packet by the source node to the destination node by using the accessed address;

means for receiving the first packet by the destination node;

means for updating the accessed address to reflect a new address of the destination node responsive to a change in the address of the destination node to the new address;

means for sending a second packet by the source node to the destination node by using the new address; and

means for receiving the second packet by the destination node.

22. A computer-readable medium containing instructions for controlling a data processing system to perform a method, the method for communicating in a network with a source node and a destination node, wherein each node has an address, the method comprising the steps of:

sending a first packet from the source node to the destination node by using the address of the destination node, wherein the source node and the destination node are programs;

receiving a new address by the source node to supersede the address of the destination node responsive to a change in the address of the destination node to the new address; and

5 sending a second packet by the source node to the destination node by using the new address.

23. The computer-readable medium of claim 22 wherein the network is a private network running on a public network infrastructure.

10 24. The computer-readable medium of claim 22 wherein the change in the address of the destination node to the new address is caused by a device on which the destination node runs physically changing locations.

COUNSEL FOR APPELLANT

LAW OFFICES
FINNEMAN, HENDERSON,
FARABOW, GARRETT,
& DUNNER, L.L.P.
1300 I STREET, N.W.
WASHINGTON, D.C. 20005
202-408-4000

ABSTRACT

Methods and systems consistent with the present invention provide a Supernet, a private network constructed out of components from a public-network infrastructure. The Supernet provides flexible and dynamic mobility support. When a destination node moves to a new location, it automatically updates the sending nodes with its new IP address. The destination node can choose among a number of ways to update the sending nodes, providing flexibility not found in conventional networks. Thus, a node can change locations repeatedly and continue to communicate directly with other nodes without the use of a proxy or other middleman.

BENEDICT & CO., INC.

LAW OFFICES
FINNEMAN, HENDERSON,
FARABOW, GARRETT,
& DUNNER, L.L.P.
1300 I STREET, N.W.
WASHINGTON, D.C. 20005
202-408-4000






Fig. 1

© 2008 All Rights Reserved
(Prior Art) © 2008 All Rights Reserved

**Fig. 2A
(Prior Art)**



**Fig. 2B
(Prior Art)**





Fig. 3

Fig. 4






Fig. 5 Diagram illustrating the architecture of two nodes (Node A and Node B) connected via a network interface.




Fig. 6A




Fig. 6B




Fig. 7




Fig. 8




Fig. 9




FIG. 10




FIG. 11